QL

PCT / IB 0 0 / 0 0 8 4 1

31. 07. 00   31. 07. 00

INVESTOR IN PEOPLE

10-009250

The Patent Office

Concept House

Cardiff Road

Newport

South Wales

NP10 8QQ

REC'D 17 JUL 2000

WIPO PCT

PRIORITY DOCUMENT

SUBMITTED OR TRANSMITTED IN COMPLIANCE WITH RULE 17.1(a) OR (b)

IB00/00 841

4

I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

I also certify that by virtue of an assignment registered under the Patents Act 1977, the application is now proceeding in the name as substituted.

I also certify that the attached copy of the request for grant of a Patent (Form 1/77) bears an amendment, effected by this office, following a request by the applicant and agreed to by the Comptroller-General.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.

Signed A Evans.

Dated 07 July 2000

INVESTOR IN PEOPLE

GB9913239.1

By virtue of a direction given under Section 30 of the Patents Act 1977, the application is proceeding in the name of

MARCONI COMMUNICATIONS SPA
Via A Negrone
1A 16153
Genova
Italy

Incorporated in Italy

[ADP No. 07923576001]

Patents Form 1/77

Patents Act 1977
(Rule 16)

The Patent Office

1/77

THE PATENT OFFICE
A
-8 JUN 1999
RECEIVED BY FAX

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
Gwent NP9 1RH

1. Your reference — P/61813.GBA/GPT U18

2. Patent application number *(The Patent Office will fill in this part)* — 9913239.1

SECTION 30 (1977 ACT) APPLICATION FILED 31.5.00

08JUN99 E452884-1 D00314
P01/7700 0.00 - 9913239.1

3. Full name, address and postcode of the or of each applicant *(underline all surnames)*

MARCONI COMMUNICATIONS LIMITED
P.O. BOX 53
NEW CENTURY PARK
COVENTRY CV3 1HJ

Patents ADP number *(if you know it)* — 7519200001

If the applicant is a corporate body, give the country/state of its incorporation — UNITED KINGDOM

4. Title of the invention

COMMUNICATIONS ARRANGEMENT

5. Name of your agent *(if you have one)*

"Address for service" in the United Kingdom to which all correspondence should be sent *(including the postcode)*

~~COLIN F. HOSTE~~
~~GEC PATENT DEPT.~~
~~WATERHOUSE LANE~~
~~CHELMSFORD~~
~~ESSEX CM1 2QX~~
U.K.

Marconi Intellectual Property
Waterhouse Lane
CHELMSFORD
Essex
CM1 2QX

Global change 3rd 7/00

Patents ADP number *(if you know it)* — 4052700

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and *(if you know it)* the or each application number

Country	Priority application number *(if you know it)*	Date of filing *(day / month / year)*
—	—	—

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application	Date of filing *(day / month / year)*
— —	—

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? *(Answer 'Yes' if:*
a) any applicant named in part 3 is not an inventor, or
b) there is an inventor who is not named as an applicant, or
c) any named applicant is a corporate body.
See note (d))

Yes.

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form	—
Description	17 ✓
Claim*(s)*	4 ✓
Abstract	1 ✓
Drawing*(s)*	6 ✓

10. If you are also filing any of the following, state how many against each item.

Priority documents	—
Translations of priority documents	—
Statement of inventorship and right to grant of a patent *(Patents Form 7/77)*	—
Request for preliminary examination and search *(Patents Form 9/77)*	1
Request for substantive examination *(Patents Form 10/77)*	—
Any other documents *(please specify)*	—

11. I/We request the grant of a patent on the basis of this application.

Signature C. F. Hoste (COLIN HOSTE) Date 8 JUNE 1999

12. Name and daytime telephone number of person to contact in the United Kingdom PETER WEBB 01245 275126

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

a) *If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.*

b) *Write your answers in capital letters using black ink or you may type them.*

c) *If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.*

d) *If you have answered 'Yes' Patents Form 7/77 will need to be filed.*

e) *Once you have filled in the form you must remember to sign and date it.*

f) *For details of the fee and ways to pay please contact the Patent Office.*

COMMUNICATIONS ARRANGEMENT

The invention relates to a communications arrangement forming part of a communications system, in particular, but not exclusively, an SDH-DCN communications system.

A typical OSI (Open Systems Interconnection) routeing scheme involving the so-called "IS-IS Routeing Protocol" is illustrated in Figure 1. In Figure 1 a wide-area network (WAN) is shown divided into two domains, each domain being split into two areas. Each area contains a number of systems, which are designated either as end-systems (ES's) or intermediate systems (IS's). The ES's, which may represent hosts or various devices (e.g. servers), may be linked to one or more IS's via either point-to-point or broadcast circuits in a LAN (Local Area Network) or, for a geographically larger area, a MAN (Metropolitan Area Network) or WAN (Wide Area Network).

Routeing of message packets from any ES in one area to another ES in the same or another area is conventionally carried out under separate routeing protocols which correspond to a particular routeing hierarchy. Routeing between ES's and IS's is by way of the ES-IS protocol; that between any two IS's within the same area is via the intra-domain IS-IS protocol (Level 1), and that between two IS's in different areas is via the intra-domain IS-IS protocol (Level 2). Routeing between two different domains is outside the scope of the IS-IS protocol. However, the protocol provides a way to disseminate the inter-domain routeing information to all the inter-area routers, or level 2 Intermediate Systems, as they are called.

Two types of routeing have traditionally been employed: static routeing and dynamic routeing. With static routeing, some Intermediate Systems in a domain store

routeing criteria of various types. Such criteria are manually entered by the operator and are used to match the destination address of a packet against the criteria, to ascertain whether the packet may be routed on the circuit to which the static route is associated. With dynamic routeing, each system keeps a table containing the state of all routes within its scope. The table is updated on a continual basis. Since dynamic routeing is adaptive, being able to take account of broken links between systems or to take account of systems themselves being out-of-service, and is also decentralized, it has clear advantages over static routeing and is therefore the dominant form of routeing currently employed, at least under the intra-domain IS-IS protocol.

The IS's are divided into two main types: level 1 (L1), which routes packets within a particular area, and level 2 (L2), which routes packets between areas and between domains. Usually, an L2 IS also has an L1 routeing function, and is therefore actually an L1/L2 IS.

For dynamic routeing to occur, the following conditions must be satisfied.

- Each IS must be apprised of the state of its neighbour ES's. In the same way, each ES must be apprised of the state of its neighbour IS's.
- Each L1 IS must be apprised of the topology of the area of which it is part.
- Each L2 IS must be apprised of the topology of the level 2 sub-domain of which it is part (that is, the partition of the domain made up of L2 IS's and of the links between them).

In order to accomplish this, all the End and Intermediate Systems in the domain exchange "hello" packets, to know who are their neighbours. IS's also generate and flood LSP's (the LSP, or Link State Protocol data unit, is a packet containing the list of the

neighbours of the originating IS), so that they become aware of the topology of the partition of the network within their scope (where the scope is the area, for level 1 IS's, and the level 2 sub-domain, for level 2 IS's). In this way, IS's are enabled to make the appropriate routeing decisions both at the L1 and at the L2 routeing level.

The dynamic link-state updating process just described does not occur on the inter-domain level, however, but instead a static routeing method has to be employed in order to route packets from one domain to another. In order to achieve this, the routeing tables of the L2 IS's are provided with "reachable address prefixes" (RAP's), which are generated either manually, or by means of a dynamic inter-domain routeing protocol. Such RAP's provide routeing criteria for the packets that may not be routed on the basis of the dynamic routeing information available (as they are addressed outside the domain). The criterion is that if the destination address of the packet begins with a pattern matching an existing prefix, it may be forwarded on the circuit associated with such a prefix (which circuit will turn out to be a domain boundary).

In an actual routeing exercise, an L1 IS will receive a packet from one of its associated ES's (note that if a system acts both as an End and as an Intermediate System, this is represented by the IS having itself as an ES neighbour). If the packet is destined for an ES in the same area, it will be routed by that L1 IS either to the destination ES directly, or via one or more other L1 IS's. If the packet is destined for an ES outside the source area, the L1 IS will pass the packet on to the nearest L2 (or L1/L2) IS in the source area (possibly passing through one or more other IS's). Once the packet gets to the L2 sub-domain, it will be passed on to an appropriate L2 (or L1/L2) IS in the destination area.

Finally, the packet is delivered by L1 routeing to the destination ES, either directly, or via one or more other L1 IS's.

In one particular type of telecommunications system, namely the SDH (Synchronous Digital Hierarchy) system, a ring arrangement of systems (called "network elements" in SDH terminology) is often employed. This is illustrated in Figure 2, in which an ethernet LAN 2 is connected to a ring 10 of network elements (NE's) A, B, C and D (only four are shown for the sake of simplicity), one of which - NE A - is designated as the "gateway NE" (GNE). The GNE (also known as the "head of the ring") is the NE which provides access to the other NE's in the ring for the Element Manager (EM). The EM is a system (normally running on a computer) which performs administrative operation on SDH NE's, such as configuration, alarm and performance data management. The DCN (Digital Communications Network) is the network that provides the support for the dialogue between the EM and the NE's. In practice, there may be tens of NE's on a ring and many tens of rings connected to a single GNE. As most of these NE's need to be functioning as IS's (because they have to route packets towards the further NE's), usually all the NE's are configured to act as IS's. Also, in practice there may be a number of GNE's present on LAN 2 (Figure 2 shows a second ring 20 with its own GNE, GNE 2), and some of the NE's in a ring may in turn have their own sub-rings.

Each NE in each ring is equipped with a couple of data communication channels (DCC's) through which it communicates with the next and the previous NE in the ring. These DCC's are shown in Figure 2 as channels "1" and "2" associated with NE's A, B, C and D.

The element manager 11, which is connected to another LAN, LAN 1, communicates with the rings 10 and 20 via the Data Communications Network (DCN) 12, via a router 13 upstream of the DCN, and via a router 14 downstream of the DCN and connected to the LAN 2. The routers are effectively IS's.

In the normal configuration, in which IS-IS dynamic routing (described earlier) is employed, router 14, all the GNE's on LAN 2 and all the NE's reachable through these GNE's (including NE's B, C and D) are located in the same IS-IS area. Since there may result a large number of IS's in that area, problems in routeing may be caused due to the restrictions in the number of IS's which the IS-IS protocol, by its design, can handle. In practice, the protocol suggests that a typical maximum configuration domain will contain at most 400 L2 IS's and at most 100 L1 IS's per area, while the domain is allowed to comprise up to 4000 systems.

The above-mentioned restrictions are due to the fact that each of the NE's in the area has, in the conventional arrangement, a complete view of the topology of the area (as explained earlier). Thus, an NE reachable through a GNE on one ring (e.g. GNE 1) has to process all the LSP's generated by any other NE on any ring in the same area (e.g. one of the NE's in ring 20), and this can lead to various problems, such as memory exhaustion, CPU overload and traffic bursts due to the routeing messages (such bursts may be particularly critical when there is a sudden change in the network topology).

The present invention aims to obviate this restriction while still employing the standard IS-IS protocol. It does this by recognising that it is not important for the rest of the network to know exactly the topology of each ring. It is only important that the DCN knows on which DCC's an NE may be reached. It is also desirable, in a real system, that

a failure on one of the links of a ring (e.g. the link between NE B and NE C in ring 10) should not make an NE impossible to reach, as long as the other route to that NE is still operative.

In accordance with the invention, there is provided a communications arrangement forming part of an SDH-DCN communications system, the arrangement comprising a network (LAN 1), a gateway network element (NE A) connected to the network and one or more further network elements (NE B, C, D) which, together with the gateway element, form at least a part of a routeing area, the gateway element acting as an interface between the further elements and the network, wherein the further elements are intermediate systems, but the gateway element and the further elements are configured such as to make the further elements appear as end systems as far as the rest of the communications system is concerned.

Preferred features and various realisations of the invention are contained in the subclaims.

Realisations of the invention will now be described with the aid of the drawings, of which:

Figure 1 is a diagram illustrating the IS-IS routing hierarchy in a typical communications configuration;

Figure 2 shows an SDH-DCN ring arrangement;

Figure 3 shows how the ring 10 of Figure 2 appears to the rest of the DCN, after the GNE of that ring has been reconfigured to enable the alternative IS-IS routeing defined by the invention to be employed;

Figures 4 to 6 illustrate alternative realisations of the invention in the case where a single GNE is involved, and

Figures 7 and 8 illustrate alternative realisations of the invention in the case where multiple GNE's are involved.

A first realisation of the invention is now described, in which the gateway NE for ring 10, i.e. NE A, is reconfigured such as to make the ring appear as a number of ES's connected to the GNE, as far as the rest of the DCN is concerned. To effect the re-configuration, NE A has its "external domain" flag manually set to TRUE for both the point-to-point circuits associated with DCC 1 and DCC 2. This setting of the flag, (which is provided by the protocol normally to set a domain boundary) in this case allows the prevention of IS-IS routeing towards the rest of ring 10. In addition, manual end-system adjacencies for NE B, C and D are entered by the operator both on DCC 1 and on DCC 2 of NE A, so that the rest of the network "sees" NE's B, C and D as ES's, reachable through DCC 1 or 2. A manual ES adjacency is a static route which may be entered manually on an IS to declare that one or more ES's in its same area are neighbours of the IS on a given circuit. The result is that the topology illustrated in Figure 3, in which NE's B, C and D appear as ES's directly connected to DCC's 1 and 2 of GNE A, is made known to the rest of the system by LSP's generated by the GNE.

Since NE A's "external domain" flag is now set on DCC 1 and 2, no IS-IS LSP's are forwarded onto the ring (though ES-IS LSP's can still flow the network between the GNE and its neighbours – but such packets are local to the circuit on which they are generated). The result is that, in this arrangement, there is no overhead due to the processing by the GNE of LSP's generated on the ring 10 (or on any other ring for which,

in practice, it might act as a gateway) and no overhead due to the processing, on the NEs of one ring (e.g. ring 10), of LSP's generated on another ring (e.g. ring 20). There are also no LSP bursts (the previously mentioned "flooding") involving the ring 10.

Although NE A has been suitably reconfigured to make ring 10 effectively a separate domain, this is not in itself sufficient to create a working system, since the other NE's on the ring, NE's B, C and D, are still set up as normal IS's which rely on the receipt of dynamic routeing information from the rest of the network, and there is now no exchange of such dynamic routeing information involving this ring. In order to make it possible for NE B, for example, to send packets out of the ring 10, it is (manually) configured as a level 2 (L2) IS having its "external domain" attribute flag set to TRUE on its own DCC 1 circuit. In addition to this, it is provided with a length-zero RAP (reachable address prefix) on that same DCC. A length-zero prefix is a RAP that matches any destination address. Thus a packet with any destination address which is to be sent out of the ring from NE B will automatically be routed through to the gateway, NE A, and out to the rest of the network.

A similar process is applied to NE C, i.e. it has its "external domain" attribute set to TRUE and it is also provided with a length-zero prefix, but both this time on circuit DCC 2, since this is the circuit which is immediately connected to NE A.

As far as NE D (and any other NE possibly present on the ring) is concerned, this does not require to be manually reconfigured at all, but can be maintained as a normally configured L1 IS with no manual routeing information (i.e. no RAP's and no manual adjacencies). In practice, NE D will route all its outgoing traffic towards either NE B or NE C, as these NE's each have an outward route by virtue of their RAP setting. NE B or

C, in turn, forward the packets they receive towards NE A. Once on NE A, all the packets are delivered to their destination via normal IS-IS dynamic routing, since NE A has a dynamically derived knowledge of the DCN topology before it.

Concerning incoming traffic entering the ring, NE A delivers any packet addressed to one of the NE's for which it has manual adjacencies configured, either on DCC 1 or on DCC 2. The rest of the network becomes aware of these adjacencies because they are reported in the level 1 (intra-area, i.e. L1) LSP's communicated from NE's B and C to the gateway and passed on from there to the rest of the system. Hence the rest of the network knows to send packets having a ring-10 address to gateway NE A.

It should be noted, as an explanatory point, that, although the NE's in the ring appear to the rest of the DCN as ES's, inside the ring they still appear to each other as a ring-configured series of IS's within a separate "domain". Hence normal dynamic routing takes place within the ring, but not between the ring and the rest of the system.

What has so far been described is a basic realisation of the invention. In reality, however, this novel approach presents a number of drawbacks. These are:

(1) It is necessary to set up two manual end-system adjacencies on NE A for each NE on the ring. Since there may be many elements on the ring, this represents an undesirable overhead in terms of human technical intervention.

(2) As manual adjacencies are static, there is no dynamic recovery against the failure of a link on the ring. Indeed, it is the strength of dynamic routeing that such failure can be circumvented. There may, however, be some recovery for those links adjacent to the gateway. For such circuits, it is

possible to monitor the Data Link connection, assuming that AITS (Acknowledged Information Transfer Service) is the transfer service mode allowed on the Data Link Layer., In this case, the Data Link layer turns out to be connection-oriented. Under these circumstances a loss of connection will be promptly signalled to the Network Layer and therefore to the gateway. (The Data Link Layer and the Network Layer are Layers 2 and 3 of the OSI system).

(3) Under normal intra-domain IS-IS routeing, a packet originated within one part of an area and bearing an address within another part of the same area should be routed to the destination area (by L1 IS routeing). In the present case, however, a message generated within ring 10 and addressed to an NE outside the ring, but within the same area (e.g. within the ring 20 in Figure 3), may be discarded by the originating NE. This is because ring 10 acts as a separate domain outside the scope of L1 routeing.

A solution to these drawbacks is now addressed.

Manual End-System Adjacency Overhead

The proposed solution here is to allow a user to enter onto the gateway NE a static route record defining a manual adjacency covering one or more ranges of consecutive System Identifiers. Thus, if it is possible to define a static route record allowing for up to, say, three ranges of up to, say, 255 System Identifiers, the user could enter with only a single data entry static routes for up to 765 NE's which may be reached via a given DCC. This solution implies that the Network Addressing Plan is designed in such a way as to have consecutive System Identifiers for the NE's belonging to the same ring. (There may,

however, be some gaps in the series). In practice this is not a real limitation on the Network Addressing Plan, since this "consecutiveness" constraint is already normally satisfied anyway.

Static Routeing Recovery Mechanism

Under normal routeing practice, whether GNE A sends an incoming packet to, say, NE D on its DCC 1 circuit or its DCC 2 circuit will depend on the comparative metrics of the two routes. The metric is a measure of the cost of establishing a route over a particular circuit according to a particular criterion chosen. Possible criteria are: the circuit's capacity to handle traffic (this is the normal "default" metric used), the transit delay of the circuit, the monetary cost of utilising the circuit, or the residual error probability of the circuit. Whichever metric is used, the cheapest circuit is chosen. In the proposed, static, configuration, the choice made by GNE A between DCC 1 and 2 is a random one, as the two DCC's each happen to have a Manual Adjacency to the destination with the same metric cost. Under normal dynamic routeing, if one of two possible circuits were "down" (not functioning), this would be dynamically communicated to the NE concerned so that it would then choose the other, regardless of whether it was the "best" (metrically speaking) route. Under static routeing this is not possible, with the result that the NE will try to send a packet along the "best" route anyway, even if it is "down".

Now, if the link between A and B, or between A and C, fails, the GNE is notified of this failure by the Data Link layer (supposing that the layer two protocol is connection oriented, which is necessary for recovery to work). In this case, the manual adjacency on the failed circuit becomes non-operational, so that the other circuit is automatically selected. But if the circuit between B and D or between C and D fails, the GNE has no way

of knowing this (as there is no longer any dynamic knowledge of the ring topology). In this case, if, for example, the link between B and D fails, and A sends to B a packet addressed to D, B sends the packet back to A. In fact, the only possible route for the packet, on B, is the length-zero prefix towards A. If A chooses to go on sending the packet to B, it enters a routeing loop until its lifetime expires and it is discarded.

The solution proposed here is the following upgrade to the implementation of the forwarding process on the NE. Suppose that the following conditions apply:

1. The NE has to forward a packet using a manual adjacency.
2. There are two manual adjacencies, say ADJ-1 and ADJ-2, which match the destination address of the packet.
3. ADJ-1 and ADJ-2 have the same metric cost.
4. The packet was received on the circuit associated with ADJ-1 (ADJ-2).

In this situation, the NE forwards the packet on ADJ-2 (ADJ-1), so that, if it is sent onto the "wrong" DCC, the first time (in the above example, towards B), when it comes back it is sent onto the right one (in this case, towards C).

<u>Intra-Area Routeing Protection</u>

If a packet bears an address in the area in which it was originated, but outside the ring, it is arranged for the packet not to be discarded by level 1 routeing if an "attached" Level 2 router (that is, an L2 IS which declares itself to have access to other areas - in this case, one of the NE's connected to the GNE) can be reached by the L1 IS handling the packet.

This improvement applies to the small DCN, where the EM and all the NE's fit in the same area, and to all DCN's in which it is necessary to route packets from one ring to another.

Although the invention has been explained largely with reference to a simple ring configuration, as illustrated in Figure 1, the invention is not limited to only such configurations. Indeed, it is not restricted to ring topologies at all.

Figures 4 to 9 depict other realisations of the invention in an SDH-DCN setting. In Figure 4a, an arrangement using exclusively IS-IS dynamic routeing is shown in which, as before, an Element Manager 11 is connected to an ethernet LAN (LAN 1) and thence via routers 13, 14 and a DCN 12 to a second LAN (LAN 2). In this case, however, a gateway NE, GNE A, has connected to its DCC 1 a network comprising IS's B, C, D and E - elements C, D and E forming a conventional ring - and has also connected to its DCC's 2 and 3 a ring comprising IS's F, G, H and I, element I itself having connected to it a ring consisting of IS's L and M.

The present invention takes the given IS configuration and splits off the IS's on DCC 1 as one group of ES's reachable on this DCC and the IS's on DCC's 2 and 3 as another group of ES's. A manual ES adjacency to systems B, C, D and E is entered by the operator on the circuit associated with DCC 1, for which the "external domain" attribute is also set. Two more manual ES adjacencies to systems F, G, H, I, L and M are entered, on the circuits associated with DCC's 2 and 3. In addition, elements B, F and G are designated as L2 IS's, and are provided with length-zero RAP's on their circuits connecting to the GNE, for which circuits the "external domain" attribute is set TRUE as

well. The topology displayed by the GNE to the rest of the DCN, by means of its own LSP's, is shown in Figure 4b.

In this case, it should be noted that, as far as the part of DCN reached on DCC 1 is concerned, there is only one boundary, since only one DCC is involved. This means that there is no "redundancy" and therefore a break in any link cannot be obviated, unlike the case where there is a failure of a link or system reachable through DCC 2/3.

Figure 5a shows a configuration in which GNE's are reached by the DCN not via ethernet, but by means of DCC directly. In the setup shown, an ADM-4 arrangement comprises NE's B and C, which act as non-gateway NE's for an STM-4 ring and at the same as GNE's for a number of STM-1 sub-rings 24, 25 and 26. It is assumed that dynamic routeing is to be maintained for the STM-4 ring, but not for the others. The invention is here brought into play to arrive at Figure 5b. In Figure 5b each of the STM-1 sub-rings 24, 25 and 26 is converted as previously described by making elements B and C "local" GNE's for their respective rings. The "external domain" attribute is set TRUE for DCC's 1 and 2 on B, (and DCC's 3 and 4 on C) and ES adjacencies are entered by the operator on B and C for the non-gateway NE's. These adjacencies are towards NE's I and J in ring 24, D and E in ring 25 and F, G and H in ring 26. As before, all these same elements (except element G) have length-zero prefixes set up on the circuits connecting them to their respective GNE's, for which the "external domain" flag is also set TRUE, and are all configured as L2 IS's, except G. NE G is not reconfigured, but is maintained as an L1 NE.

A "bus" type of topology such as that illustrated in Figure 6a will, according to the invention, be transformed into the configuration shown in Figure 6b. Here a manual end-

system adjacency to B, C, D and E is entered on the GNE A for the DCC connecting it to B. The "external domain" attribute is also set on the circuit associated with this DCC. NE B is reconfigured to act as a L2 IS, with a length-zero prefix associated with the circuit of the DCC connecting B to A, for which circuit the "external domain" attribute is set as well. No reconfiguration is needed on the other NE's.

The invention may also be applied to a topology in which there are two GNE's, as illustrated in Figure 7. In Figure 7a GNE's A and F interface with respective LAN's 1 and 2 and form a ring with IS's B, C, D and E. In this case, manual end-system adjacencies associated with the circuits using DCC's 1 of GNE's A and F have to be entered on such GNE's for NE's B and C. In the same way, manual end-system adjacencies associated with the circuits using DCC's 2 of GNE's A and F have to be entered on these GNE's for NE's D and E. Again, the appropriate "external domain" attributes are set and length-zero RAP's supplied on B, C, D and E, which are also designated to act as L2 IS's. (See Figure 7b).

Note that manual adjacencies here are only created towards NE's which are reachable on a given DCC without crossing through another GNE. Thus in Figure 7, no adjacency to NE's B and C is created in association with DCC 2 of GNE A, since these NE's are reachable on this DCC only by crossing through GNE F.

In this particular configuration the ring is susceptible to problems resulting from a single link or other failure on the ring. To compensate for this drawback it is possible to incorporate the principles embodied in the co-pending UK patent application GB 9805247.5 of priority date 31 July 1997 and filed in the name of GPT Limited, the application being titled "M A Alternate Routeing".

Figure 8a shows a case in which three GNE's are present (A, B, C). It is assumed that these GNE's all fit in the same area that also contains routers 13, 14 and 15. The fact that the external routers all fit in the same area means that either the area is connected (that is, it is possible to go from any router to any other by following a DCN path only crossing through routers in the area), or it is partitioned, in which case, the external routers need to implement the partition repair feature of the IS-IS protocol. In this case the present invention reconfigures the topology to arrive at the scheme of Figure 8b, in which manual adjacencies are set up as will now be described:

1. GNE A: one manual adjacency on DCC 1 to NE's D, E and, F, and one on DCC 2 for NE's K and L;
2. GNE B: one manual adjacency on DCC 1 to NE's K and L, and one on DCC 2 for NE's G, H, I and J;
3. GNE C: one manual adjacency on DCC 1 to NE's G, H, I and J, and one on DCC 2 to NE's D, E and F.

"External domain" flags are set in the DCC's of the GNE's and in the circuits of NE's D, F K, L, G and J which are connected to the GNE's, while the appropriate length-zero RAP's are set up in these circuits as well.

The extension of the Figure 9 scheme to a number of GNE's more than three is easy to implement if the following condition is satisfied, namely that manual adjacencies are created on each GNE's only for the NE's comprised between it and the next closest GNE's on the same ring.

The constraint that the GNE's must be in the same area might seem restrictive at first glance, but this will often not be a problem in practice, since the number of acting IS's

according to the invention will not be great, in view of the fact that most NE's will end up looking as ES's to the rest of the DCN. Hence quite a large area can be covered without the maximum number of IS's suggested by the IS-IS protocol being exceeded.

Finally, although the invention has so far been described with reference to the SDH system, it is also applicable to other communications systems which use the standard IS-IS routeing protocol.

While reference designators have been incorporated into the claims in order to facilitate the understanding thereof, these are not to be construed as being in any way limiting to the scope of the claims.

CLAIMS

1. A communications arrangement forming part of an SDH-DCN communications system, the arrangement comprising a network (LAN 1), a gateway network element (NE A) connected to the network and one or more further network elements (NE B, C, D) which, together with the gateway element, form at least a part of a routeing area, the gateway element acting as an interface between the further elements and the network, wherein the further elements are intermediate systems, but the gateway element and the further elements are configured such as to make the further elements appear as end systems as far as the rest of the communications system is concerned.

2. Communications arrangement according to Claim 1, in which the gateway element has one or more digital communication channels (DCC's 1, 2) connected to respective ones of said further elements, each of the one or more channels being provided with manual end-system adjacencies for all the further elements and having its "external domain" attribute flag set TRUE.

3. Communications arrangement according to Claim 2, in which those further elements which are directly connected to the gateway element by a digital communication channel are configured as a Level 2 Intermediate System and have their "external domain" attribute flag set TRUE for the circuit using said digital communication channel, and in which said digital communication channel is supplied with a length-zero reachable address prefix.

4. Communications arrangement according to Claim 3, in which the gateway element has two digital communication channels each of which gives access to one of the further elements on a corresponding channel thereof, the further elements being connected in a chain configuration so to form a ring with the gateway element.

5. Communications arrangement according to Claim 3, in which the gateway element (GNE A) is connected to one of the further elements (B, C, D, E) by only one digital communication channel (DCC).

6. Communications arrangement according to Claim 5, in which there is set up on the said one gateway channel a manual end-system adjacency for all the further elements, the "external domain" attribute flag is set TRUE for that one channel and for the corresponding channel of one of the further elements which terminates the channel at the other end thereof, and a length-zero prefix is set up on the corresponding channel of the said terminating further element, said terminating further element being configured as a level 2 intermediate system.

7. Communications arrangement according to Claim 3, in which there are two gateway elements (GNE A, F) connected to respective networks (LAN 1, 2) and a plurality of further elements (B, C, D, E) connected between the two gateway elements.

8. Communications arrangement according to Claim 7, in which each of the gateway elements and further elements has two digital communication channels, a first channel

(DCC 1) of one gateway element (GNE F) being connected to a channel (DCC 1) of a first of the further elements (C), a second channel (DCC 2) of the same gateway element (GNE F) being connected to a channel (DCC 1) of a second of the further elements (E), a first channel (DCC 1) of the other gateway element (GNE A) being connected to a channel (DCC 2) of a third of the further elements (B), and a second channel (DCC 2) of the other gateway element (GNE A) being connected to a channel (DCC 2) of a fourth of the further elements (D).

9. Communications arrangement according to Claim 8, in which each channel of those further elements which are directly connected to at least one of the gateway elements has its "external domain" attribute flag set TRUE and has a reachable address prefix of length zero on the circuit connecting it to the gateway element and in which the first channel (DCC 1) of each gateway element (GNE A, F) is set with manual end-system adjacencies for the first and second of the further elements (B, C) and the second channel (DCC 2) of each gateway element (GNE A, F) is set with manual end-system adjacencies for the third and fourth of the further elements (D, E).

10. Communications arrangement according to any one of the preceding claims, in which the gateway element comprises a static route record in which has been manually entered one or more ranges of consecutive system identifiers corresponding to the manual end-system adjacencies.

11. Communications arrangement according to any one of the preceding claims, in which a change on the IS-IS forwarding process is implemented such that, if there are two, equal-cost manual adjacencies matching the destination address of a given packet and one of these is associated with the circuit on which the packet was received, the packet is forwarded onto the other circuit.

12. Communications arrangement according to any one of the preceding claims, in which a message packet, which is generated in the routeing area or partial routeing area defined by the gateway element and further elements and is destined for a network element outside that area or part-area, is not discarded by a Level 1 intermediate system handling that packet if that intermediate system has access to a Level 2 intermediate system forming part of the area or part-area.

13. Communications arrangement as claimed in any one of the preceding claims, in which the communications system of which it forms a part is an SDH communications system.

14. Communications arrangement substantially as shown in, or as hereinbefore described with reference to, Figure 3 and Figure 4, Figure 3 and Figure 5, Figure 3 and Figure 6, Figure 3 and Figure 7 or Figure 3 and Figure 8 of the drawings.

ABSTRACT

COMMUNICATIONS ARRANGEMENT

In a communications system, especially an SDH DCN system in which most network elements (NE's) function as intermediate systems (IS's) rather than end systems (ES's), it is made possible to reduce the constraints on the DCN topology posed by the IS-IS routeing protocol by configuring the gateway element (GNE) and the NE's directly connected to it in such a way that the non-gateway NE's appear as ES's to the rest of the system. This is achieved by setting up on each of the digital communication channels (DCC's) of the GNE manual end-system adjacencies for all the NE's reachable through that DCC; setting the "external domain" attribute of these DCC's to TRUE and setting the same attributes of the corresponding DCC's of those NE's which are immediate neighbours of the GNE likewise TRUE; and supplying those DCC's with length-zero reachable address prefixes (RAP's) and configuring the neighbour NE's as level 2 Intermediate Systems. The invention applies especially to SDH ring topologies, in which case the gateway element has two DCC's, but also to bus topologies, in which case the gateway has only one DCC. Systems with more than one gateway element are also catered for.

(FIGURE 3)

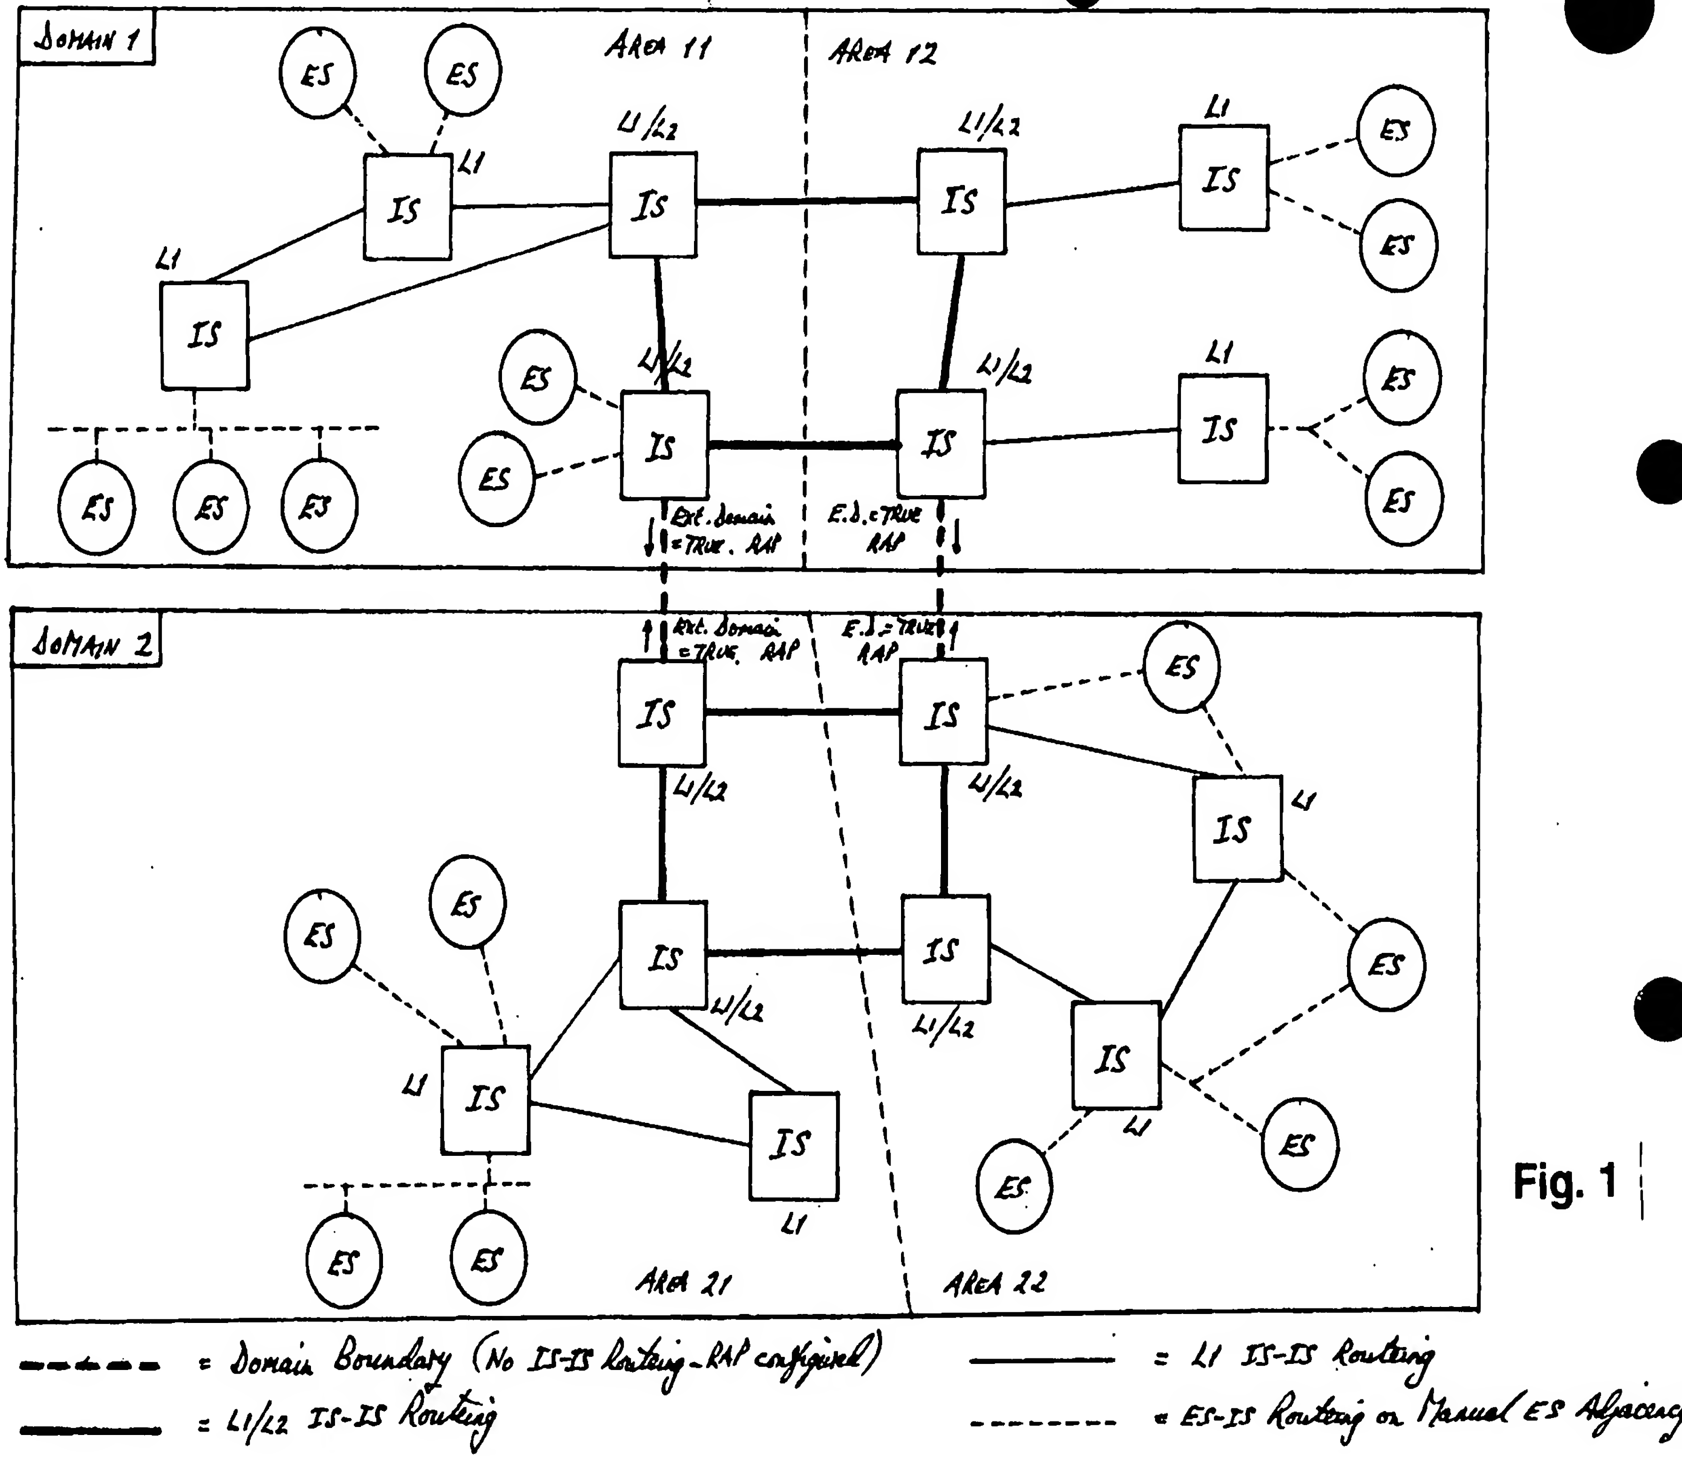
DOMAIN 1
AREA 11
AREA 12
DOMAIN 2
AREA 21
AREA 22
IS
ES
L1
L1/L2
Ext. Domain =TRUE. RAP
E.D.=TRUE RAP
Ext. Domain =TRUE, RAP
E.D.=TRUE RAP
= Domain Boundary (No IS-IS Routing - RAP configured)
= L1/L2 IS-IS Routing
= L1 IS-IS Routing
= ES-IS Routing or Manual ES Adjacency

Fig. 1

This Page Blank (uspto)

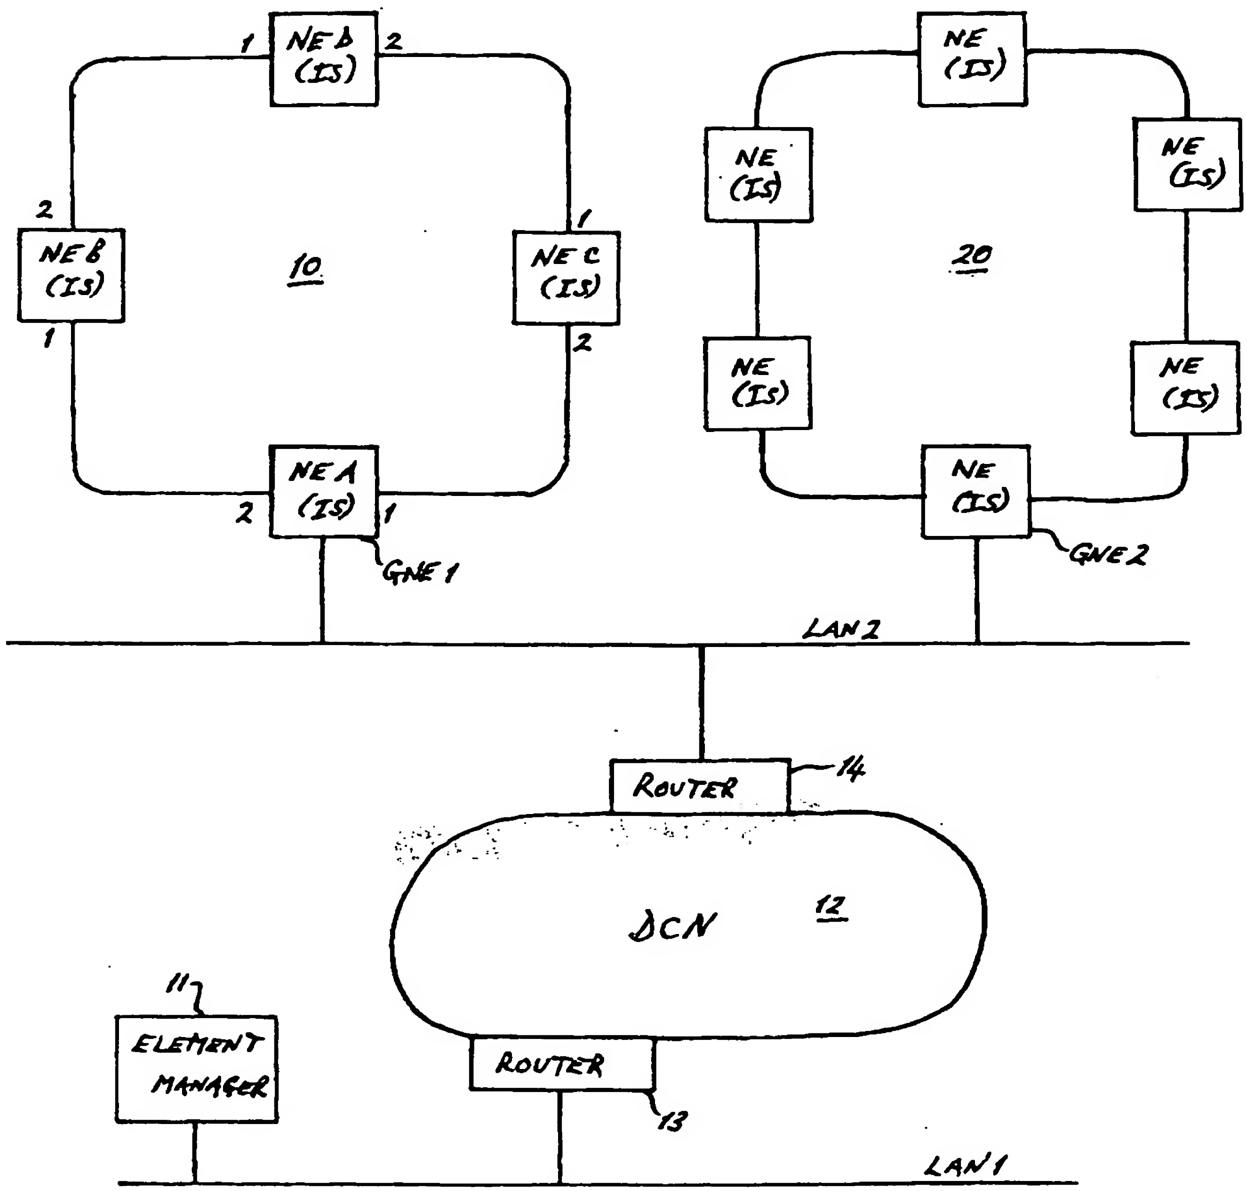
NE D
(IS)
1
2
NE B
(IS)
2
1
NE C
(IS)
1
2
10
NE A
(IS)
2
1
GNE 1
NE
(IS)
NE
(IS)
NE
(IS)
20
NE
(IS)
NE
(IS)
NE
(IS)
GNE 2
LAN 2
ROUTER
14
DCN
12
11
ELEMENT
MANAGER
ROUTER
13
LAN 1

Fig. 2

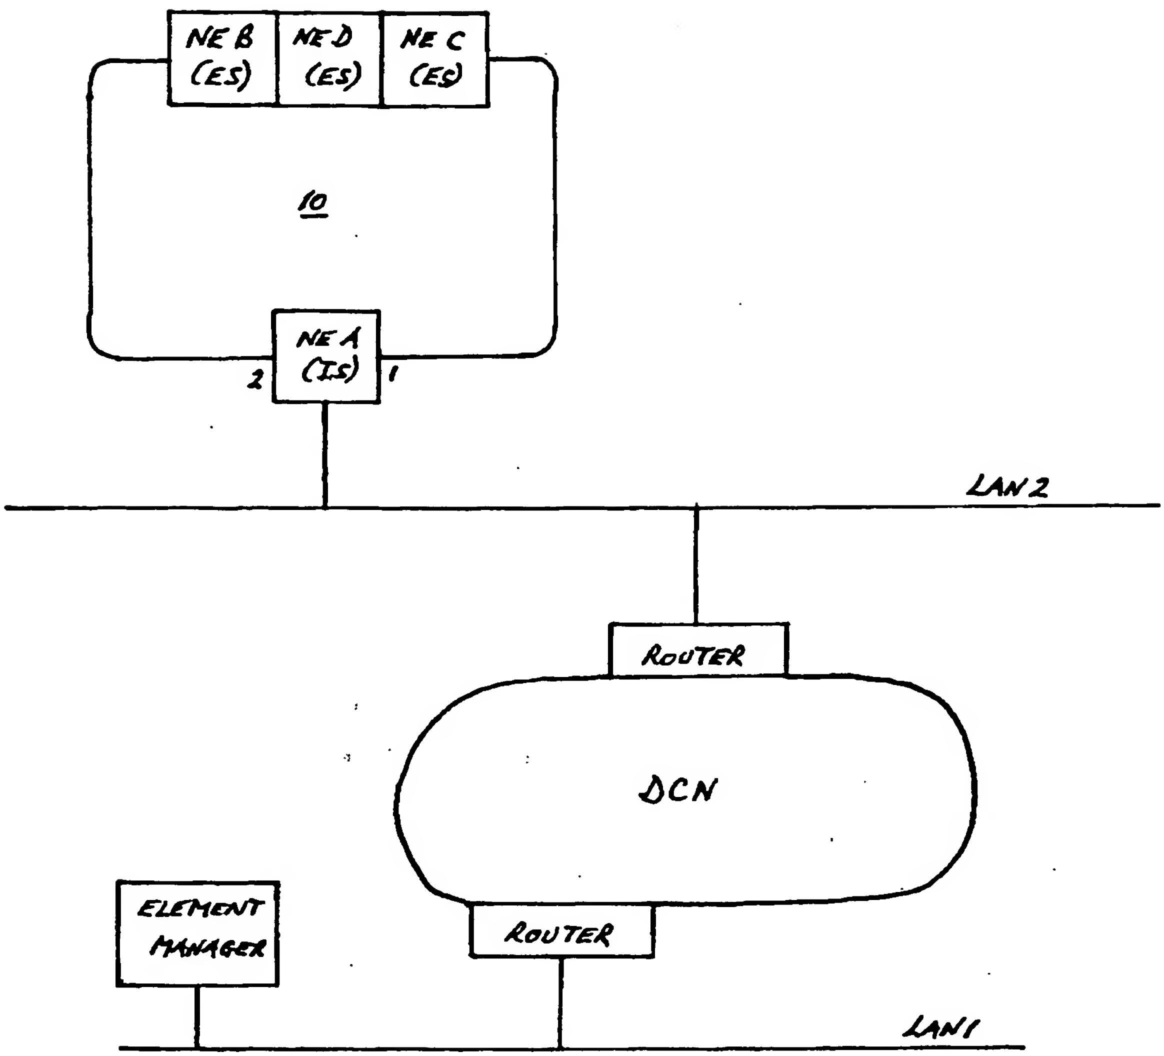
NE B (ES)
NE D (ES)
NE C (ES)
10
NE A (IS)
2
1
LAN 2
ROUTER
DCN
ELEMENT MANAGER
ROUTER
LAN1

Fig. 3

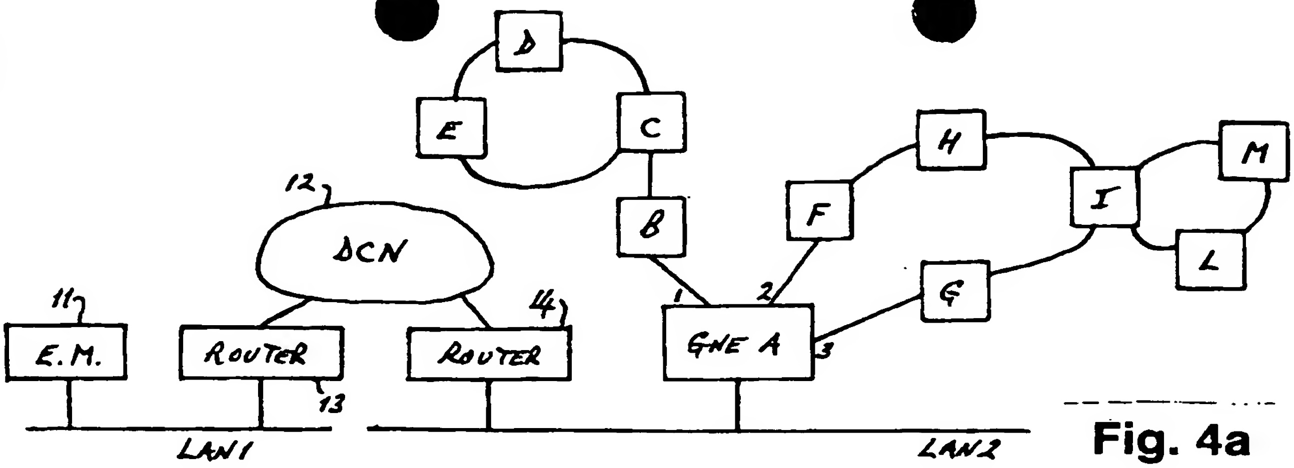

Fig. 4a

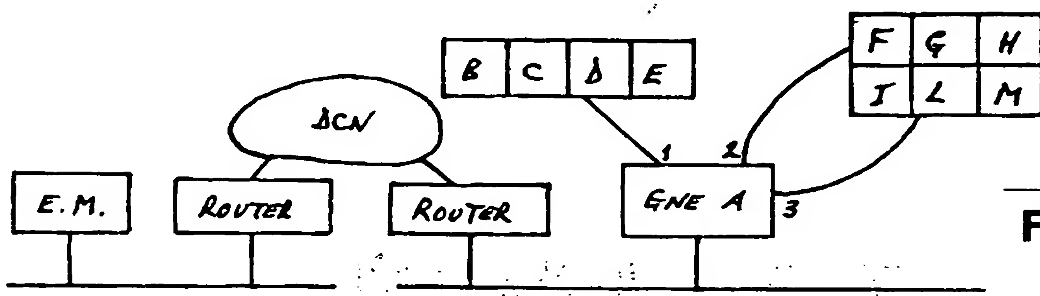

Fig. 4b

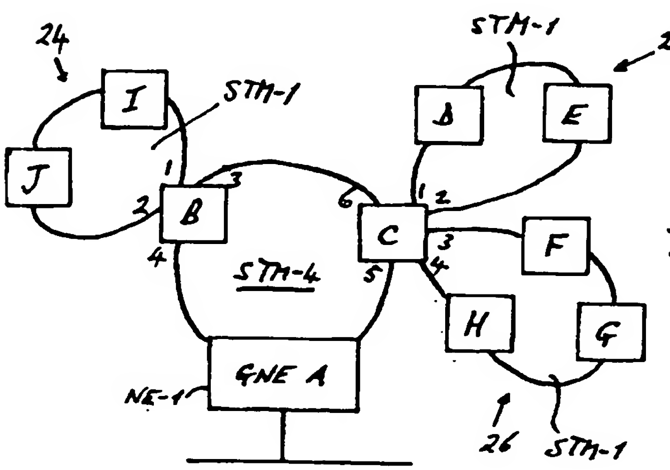

Fig. 5a

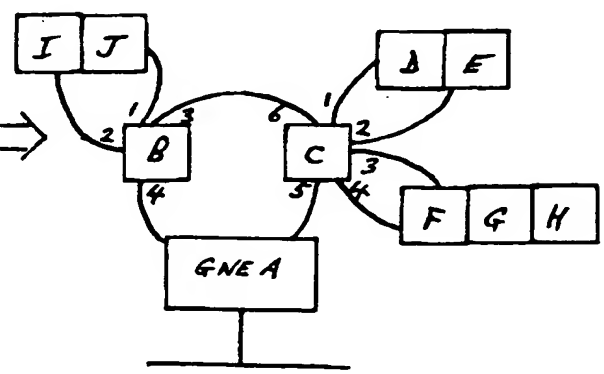

Fig. 5b

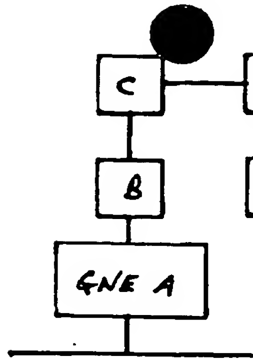

Fig. 6a

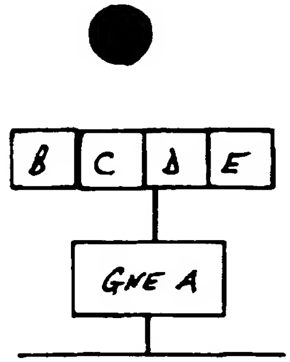

Fig. 6b

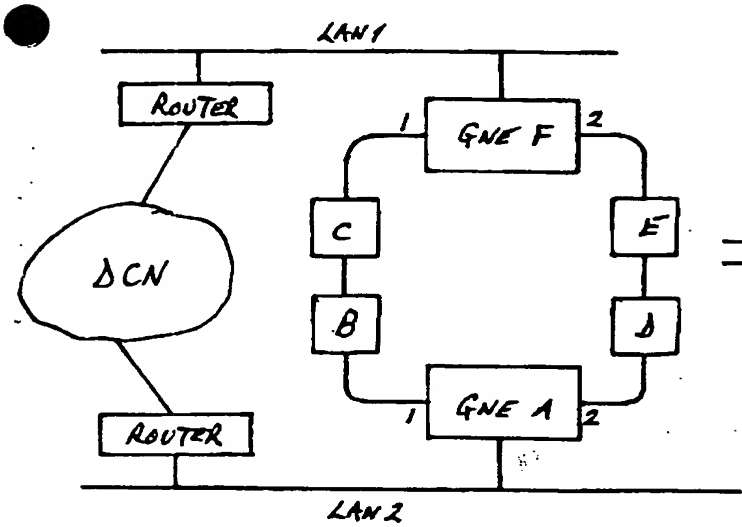

Fig. 7a

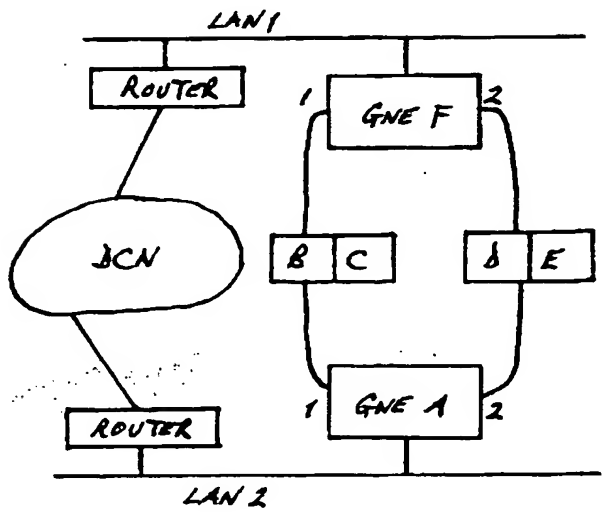

Fig. 7b

This Page Blank (uspto)

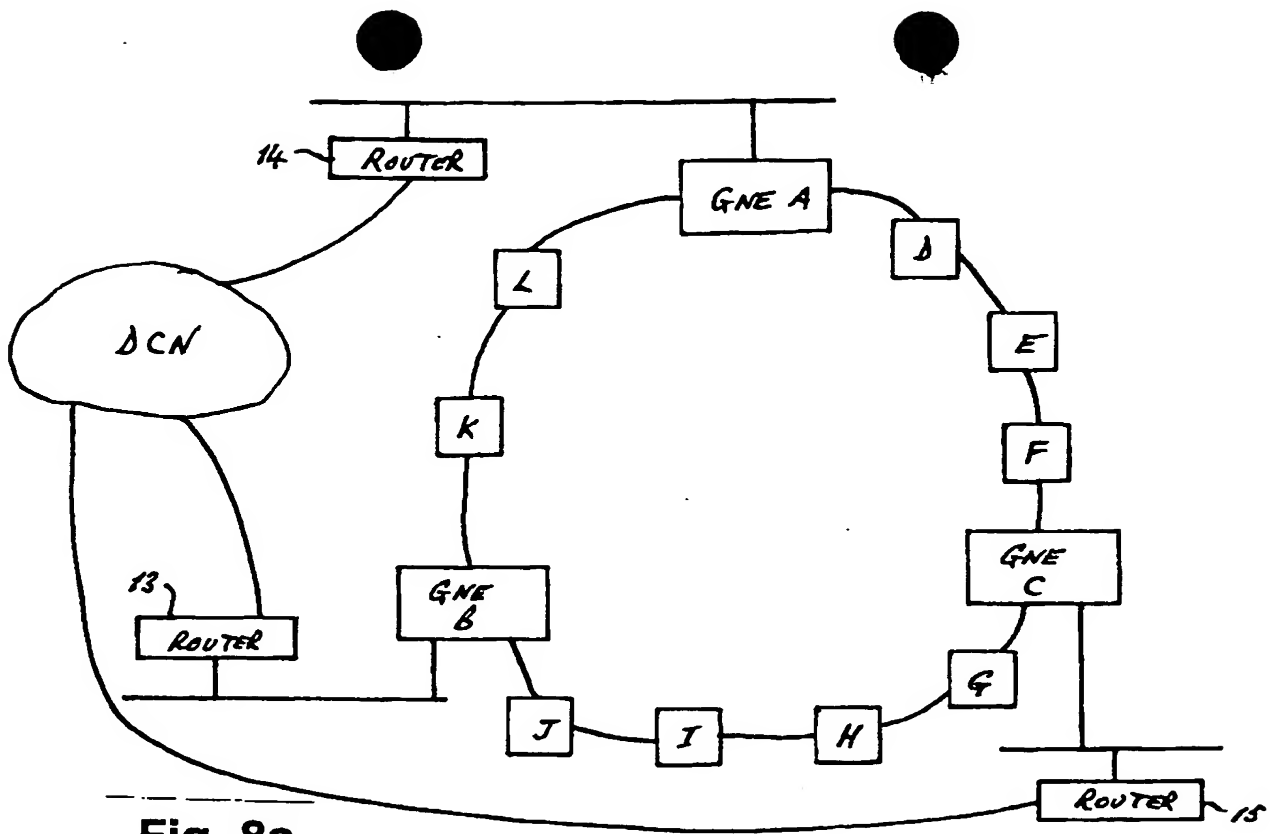

14
ROUTER
GNE A
D
L
DCN
E
K
F
13
GNE B
GNE C
ROUTER
G
J
I
H
ROUTER
15

Fig. 8a

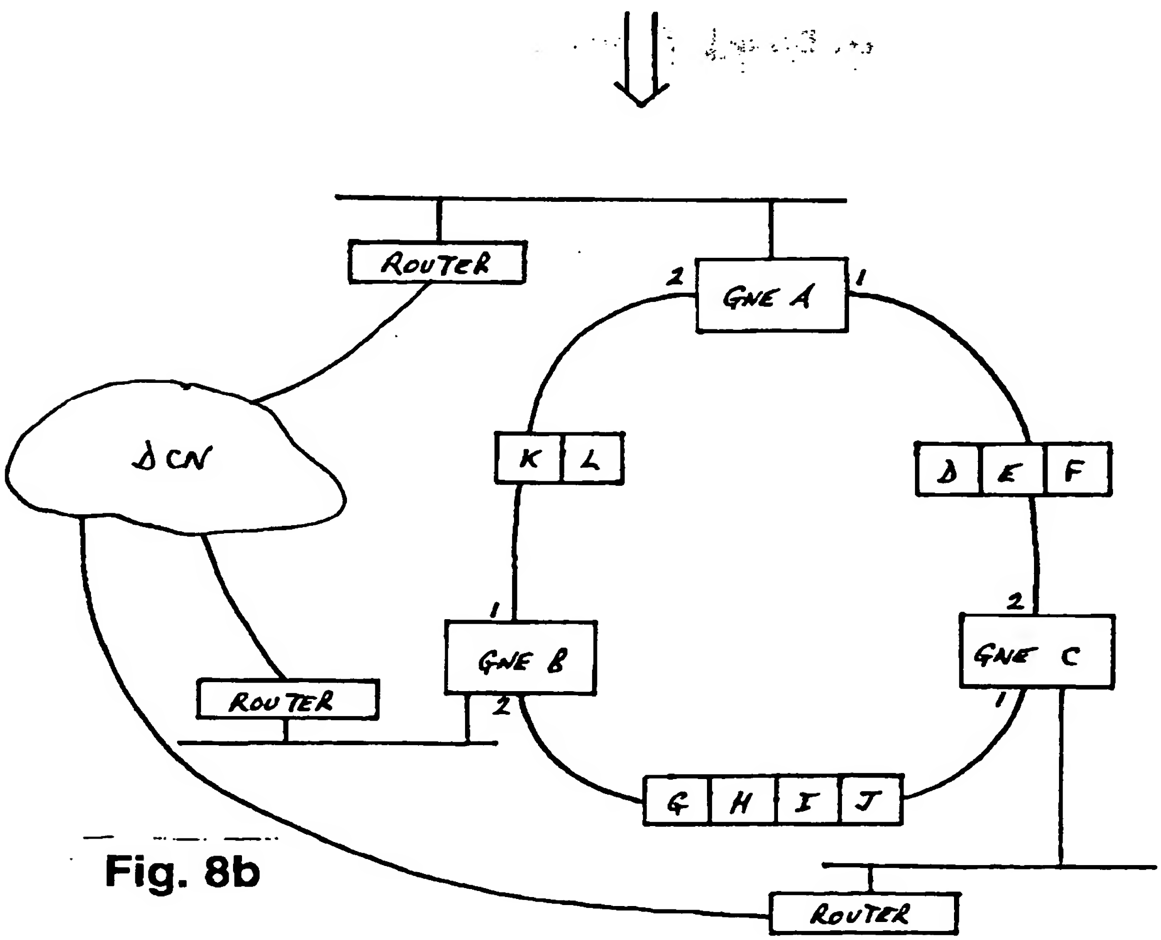

ROUTER
2
GNE A
1
DCN
K
L
D
E
F
1
2
GNE B
GNE C
ROUTER
2
1
G
H
I
J
ROUTER

Fig. 8b